

224



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/595,776	06/16/2000	Enric Musoll	P3810	1143
24739	7590	10/19/2004	EXAMINER	
CENTRAL COAST PATENT AGENCY PO BOX 187 AROMAS, CA 95004			HUISMAN, DAVID J	
			ART UNIT	PAPER NUMBER
			2183	

DATE MAILED: 10/19/2004

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No.

09/595,776

Applicant(s)

MUSOLL ET AL.

Examiner

David J. Huisman

Art Unit

2183

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 07 September 2004.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-20 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-20 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☒ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 16 June 2000 is/are: a) ☐ accepted or b) ☒ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
- ☐ Certified copies of the priority documents have been received.
 - ☐ Certified copies of the priority documents have been received in Application No. _____.
 - ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☒ Notice of References Cited (PTO-892) †
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☒ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date 07 September 2004. ‡
- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____
- 5) ☐ Notice of Informal Patent Application (PTO-152)
- 6) ☐ Other: _____

DETAILED ACTION

1. Claims 1-20 have been examined.

Papers Submitted

2. It is hereby acknowledged that the following papers have been received and placed of record in the file: Extension of Time, Information Disclosure Statement, and Amendment as received on 9/7/04.

Specification

3. The title of the invention is not descriptive. A new title is required that is clearly indicative of the invention to which the claims are directed. The examiner suggests incorporating the ideas of using a cache miss prediction to determine which thread instructions are to be fetched.

Drawings

4. The drawings are objected to under 37 CFR 1.83(a). The drawings must show every feature of the invention specified in the claims. Therefore, the specifics of claims 16, 17, and 20 must be shown or the feature(s) canceled from the claim(s). No new matter should be entered.

Corrected drawing sheets in compliance with 37 CFR 1.121(d) are required in reply to the Office action to avoid abandonment of the application. Any amended replacement drawing sheet should include all of the figures appearing on the immediate prior version of the sheet, even if only one figure is being amended. The figure or figure number of an amended drawing

Art Unit: 2183

should not be labeled as “amended.” If a drawing figure is to be canceled, the appropriate figure must be removed from the replacement sheet, and where necessary, the remaining figures must be renumbered and appropriate changes made to the brief description of the several views of the drawings for consistency. Additional replacement sheets may be necessary to show the renumbering of the remaining figures. The replacement sheet(s) should be labeled “Replacement Sheet” in the page header (as per 37 CFR 1.84(c)) so as not to obstruct any portion of the drawing figures. If the changes are not accepted by the examiner, the applicant will be notified and informed of any required corrective action in the next Office action. The objection to the drawings will not be held in abeyance.

Claim Objections

5. Claim 6 is objected to because of the following informalities: In line 1, replace “steams” with --streams--. Appropriate correction is required.

6. Claim 11 is objected to because of the following informalities: In line 9, replace “stream” with --streams--. Appropriate correction is required.

Response to Arguments

7. Applicant's arguments with respect to claims 1-13 have been considered but are moot in view of the new ground(s) of rejection, as they are necessitated by applicant's amendments.

Claim Rejections - 35 USC § 102

Art Unit: 2183

8. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

9. Claims 1-14 and 18-19 are rejected under 35 U.S.C. 102(b) as being anticipated by Yoaz et al., “Speculative Techniques for Improving Load Related Scheduling,” May 1999 (herein referred to as Yoaz). In addition, Parady, U.S. Patent No. 5,933,627 (as applied in the previous Office Action and herein referred to as Parady) is cited as extrinsic evidence for showing that it is common to have separate hardware streams for each thread.

10. Referring to claim 1, Yoaz has taught in a processor having multiple hardware streams supporting multiple data threads, and a data cache, a system for fetching instructions from a selected one of the multiple hardware streams to a pipeline, the system comprising:

a) multiple hit/miss predictors, each associated with a corresponding one of the multiple hardware streams, said each configured to forecast whether corresponding instructions from said corresponding one of the multiple hardware streams will hit or miss the data cache. See page 47, column 1, and note the paragraph beginning with “Another...”. Clearly, multiple threads exist, and a predictor is used for each thread. Consequently, there are multiple predictors.

b) a fetch algorithm, coupled to said multiple hit/miss predictors, configured to select, on a cycle-by-cycle basis, the selected one of the multiple hardware streams from which to fetch the instructions. Again, see page 47, column 1, and note the paragraph beginning with “Another...”. The fetch algorithm includes the steps of fetching from a current thread as long as a miss is not predicted, and when a miss is predicted, switching to and fetching instructions from a second

Art Unit: 2183

thread. It should be realized that this happens on a cycle-by-cycle basis, as instructions are executed every cycle. Whenever a load instruction appears, the prediction must be made, and if multiple loads occur in a row then the predictions are made for those consecutive cycles.

11. Referring to claim 2, Yoaz has taught a system as described in claim 1. Yoaz has further taught that a hit prediction precipitates no change in the fetching of the instructions. See page 47, column 1, and note the paragraph beginning with "Another...". It should be realized that Yoaz only discloses thread switching based on a predicted cache miss because that is when an expensive main memory access is likely required. When there is a hit in the cache, no main memory access is required and therefore, no switch needs to be made.

12. Referring to claim 3, Yoaz has taught a system as described in claim 1. Yoaz has further taught that a miss prediction results in switching the fetching to a different one of the multiple hardware streams. See page 47, column 1, and note the paragraph beginning with "Another...".

When a miss prediction is made, a thread switch occurs. That is, the system will begin fetching a new thread instead of sitting idle while the current thread finishes the time-consuming main memory access. Each stream is a flow of data from a source to a sink. That is from a first thread to the execution units is a first stream, a second thread to the execution units is a second stream, and so on. A hardware stream is nothing more than the hardware which supports this flow of instructions. In addition, Parady has shown that multiple hardware streams exist for multiple threads. See Fig.3, and note that the wires coupling the dispatch unit and each of the instruction buffers forms a multiple hardware streams, one for each thread, in which thread instructions may flow.

13. Referring to claim 4, Yoaz has taught a system as described in claim 1. Yoaz has further taught that said each of said multiple hit/miss predictors generates a confidence level value, and said confidence level value is used by said fetch algorithm to select the selected one of the multiple hardware streams. See page 47, column 1, and note the paragraph beginning with "Another...". Clearly, the predictors will predict a miss to occur or not to occur. This prediction is a confidence level value in that if a miss is predicted to occur, then the predictor is confident that the current thread will miss the cache, whereas if a miss is not predicted to occur, then the predictor is confident that the current thread will hit the cache. These confidence values are used to determine whether fetching will continue from the current thread or whether a thread switch occurs and fetching will begin from a new thread.

14. Referring to claim 5, Yoaz has taught a system as described in claim 1. Yoaz has further taught that a said multiple hit/miss predictors operate at a dispatch level to optimize the dispatch of consumer instructions by predicting latency of data. See page 47, column 1, and note the paragraph beginning with "Another...". It should be noted that the predictions determine which thread will be fetched from. Consequently, the predictors also determine which threads will be dispatched, as instructions which are fetched are also dispatched, otherwise they cannot be executed. And, it is known that consumer instructions are instructions which merely access some register or memory location to perform the desired operation, and therefore, it is expected that consumer instructions will be dispatched. In addition, the predictors predict the latency of data in that if a miss is predicted, then it is predicted that it will take a long latency to retrieve the data (since main memory is accessed). However, if a miss is not predicted, then it is predicted that a short latency is required as the data will be available in the cache.

15. Referring to claim 6, Yoaz has taught a processor having multiple hardware streams supporting multiple data threads, the processor comprising:

a) a data cache, comprising a plurality of levels. See page 47, column 1, Fig.3, and the paragraphs beginning with “Employing...” and “Another...”. It is clear that there is an L1 and L2 cache.

b) multiple hit/miss predictors, each associated with a corresponding one of the multiple hardware streams, said each configured to forecast whether corresponding instructions from said corresponding one of the multiple hardware streams will hit or miss said data cache (see page 47, column 1, and note the paragraph beginning with “Another...”. Clearly, multiple threads exist, and a predictor is used for each thread. Consequently, there are multiple predictors. Also, each stream is a flow of data from a source to a sink. That is, from a first thread to the execution units is a first stream, a second thread to the execution units is a second stream, and so on. A hardware stream is nothing more than the hardware which supports this flow of instructions (again, see Parady as described in the rejection of claim 1), said each of said multiple hit/miss predictors comprising:

b1) a plurality of hit/miss predictors, each configured to forecast whether said corresponding instructions from said corresponding one of the multiple hardware streams will hit or miss one or more of said levels. See page 47, column 1, and note the paragraph beginning with “Another...”. Note that it is predicted if the thread will miss the L2 cache.

b2) a fetch algorithm, coupled to said multiple hit/miss predictors, configured to select, on a cycle-by-cycle basis, the selected one of the multiple hardware streams from which

to fetch the instructions wherein said fetch algorithm selects the selected one of the multiple hardware streams based upon whether said corresponding instructions from said corresponding one of the multiple hardware streams will hit or miss said one or more of said levels. Again, see page 47, column 1, and note the paragraph beginning with “Another...”. The fetch algorithm includes the steps of fetching from a current thread as long as a miss is not predicted, and when a miss is predicted, switching to and fetching instructions from a second thread. It should be realized that this happens on a cycle-by-cycle basis, as instructions are executed every cycle. Whenever a load instruction appears, the prediction must be made, and if multiple loads occur in a row then the predictions are made for those consecutive cycles.

16. Referring to claim 7, Yoaz has taught a processor as described in claim 6. Yoaz has further taught that a hit prediction precipitates no change in the fetching of the instructions. See page 47, column 1, and note the paragraph beginning with “Another...”. It should be realized that Yoaz only discloses thread switching based on a predicted cache miss because that is when an expensive main memory access is likely required. When there is a hit in the cache, no main memory access is required and therefore, no switch needs to be made.

17. Referring to claim 8, Yoaz has taught a processor as described in claim 8. Yoaz has further taught that a miss prediction results in switching the fetching to a different one of the multiple hardware streams. See page 47, column 1, and note the paragraph beginning with “Another...”. When a miss prediction is made, a thread switch occurs. That is, the system will begin fetching a new thread instead of sitting idle while the current thread finishes the time-consuming main memory access. Each stream is a flow of data from a source to a sink. That is

Art Unit: 2183

from a first thread to the execution units is a first stream, a second thread to the execution units is a second stream, and so on. A hardware stream is nothing more than the hardware which supports this flow of instructions.

18. Referring to claim 9, Yoaz has taught a processor as described in claim 6. Yoaz has further taught that said each of said multiple hit/miss predictors generates a confidence level value, and said confidence level value is used by said fetch algorithm to select the selected one of the multiple hardware streams. See page 47, column 1, and note the paragraph beginning with "Another...". Clearly, the predictors will predict a miss to occur or not to occur. This prediction is a confidence level value in that if a miss is predicted to occur, then the predictor is confident that the current thread will miss the cache, whereas if a miss is not predicted to occur, then the predictor is confident that the current thread will hit the cache. These confidence values are used to determine whether fetching will continue from the current thread or whether a thread switch occurs and fetching will begin from a new thread.

19. Referring to claim 10, Yoaz has taught a system as described in claim 6. Yoaz has further taught that a said multiple hit/miss predictors operate at a dispatch level to optimize the dispatch of consumer instructions by predicting latency of data. See page 47, column 1, and note the paragraph beginning with "Another...". It should be noted that the predictions determine which thread will be fetched from. Consequently, the predictors also determine which threads will be dispatched, as instructions which are fetched are also dispatched, otherwise they cannot be executed. And, it is known that consumer instructions are instructions which merely access some register or memory location to perform the desired operation, and therefore, it is expected that consumer instructions will be dispatched. In addition, the predictors predict the latency of

Art Unit: 2183

data in that if a miss is predicted, then it is predicted that it will take a long latency to retrieve the data (since main memory is accessed). However, if a miss is not predicted, then it is predicted that a short latency is required as the data will be available in the cache.

20. Referring to claim 11, Yoaz has taught in a processor having multiple hardware streams supporting multiple data threads, and a data cache, a method for fetching instructions from a selected one of the multiple hardware streams to a pipeline, the method comprising:

a) for each of the multiple hardware streams, making a hit/miss prediction by a corresponding one of associated hit/miss predictors as to whether corresponding instructions for the each of the multiple hardware streams previously fetched will hit or miss the cache. See page 47, column 1, and note the paragraph beginning with “Another...”. Clearly, multiple threads exist, and a predictor is used for each thread. Consequently, there are multiple predictors.

b) selecting, on a cycle-by-cycle basis, the selected one of the multiple hardware streams from which to fetch the instructions. Again, see page 47, column 1, and note the paragraph beginning with “Another...”. The method includes the steps of fetching from a current thread as long as a miss is not predicted, and when a miss is predicted, switching to and fetching instructions from a second thread. It should be realized that this happens on a cycle-by-cycle basis, as instructions are executed every cycle. Whenever a load instruction appears, the prediction must be made, and if multiple loads occur in a row then the predictions are made for those consecutive cycles.

21. Referring to claim 12, Yoaz has taught a method as described in claim 11. Yoaz has further taught that said making comprises generating a confidence level value, and employing the confidence level value to select the selected one of the multiple hardware streams. See page 47, column 1, and note the paragraph beginning with “Another...”. Clearly, the predictors will

Art Unit: 2183

predict a miss to occur or not to occur. This prediction is a confidence level value in that if a miss is predicted to occur, then the predictor is confident that the current thread will miss the cache, whereas if a miss is not predicted to occur, then the predictor is confident that the current thread will hit the cache. These confidence values are used to determine whether fetching will continue from the current thread or whether a thread switch occurs and fetching will begin from a new thread.

22. Referring to claim 13, Yoaz has taught a method as described in claim 11. Yoaz has further taught operating the multiple hit/miss predictors at a dispatch level to optimize the dispatch of consumer instructions by predicting latency of data. See page 47, column 1, and note the paragraph beginning with "Another...". It should be noted that the predictions determine which thread will be fetched from. Consequently, the predictors also determine which threads will be dispatched, as instructions which are fetched are also dispatched, otherwise they cannot be executed. And, it is known that consumer instructions are instructions which merely access some register or memory location to perform the desired operation, and therefore, it is expected that consumer instructions will be dispatched. In addition, the predictors predict the latency of data in that if a miss is predicted, then it is predicted that it will take a long latency to retrieve the data (since main memory is accessed). However, if a miss is not predicted, then it is predicted that a short latency is required as the data will be available in the cache.

23. Referring to claim 14, Yoaz has taught a method as described in claim 11. Yoaz has not explicitly taught that the processor comprises a fine-grained multistreaming processor that concurrently executes the instructions from the multiple hardware streams. However, Official Notice is taken that fine-grained multithreading (FMT) is well known and expected in the art.

Art Unit: 2183

FMT allows expedite the completion of all of the threads being worked on, as every cycle, a next thread is switched in. Consequently, overall throughput is increased. And, it should be realized that Yoaz has taught that the prediction method will work for multithreaded systems, without explicitly saying what types (FMT, CMT, etc.). A person of ordinary skill in the art would have recognized that even with FMT, load instructions will be executed, and therefore, this prediction scheme will be useful in avoiding switching back to the thread that is predicted to miss the cache. As a result, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Yoaz to be a fine-grained multistreaming system so that the advantages of an FMT system may be achieved.

24. Referring to claim 18, Yoaz has taught a method as described in claim 11. Yoaz has further taught that said selecting comprises switching the fetching to a different one of the multiple hardware streams. See page 47, column 1, and note the paragraph beginning with "Another...".

25. Referring to claim 19, Yoaz has taught a method as described in claim 11. Yoaz has further taught that the data cache comprises a first level and a second level (see page 47, column 1, Fig. 3, and the paragraphs beginning with "Employing..." and "Another...". It is clear that there is an L1 and L2 cache.), and wherein said making comprises:

a) first forecasting whether said corresponding instructions from the corresponding one of the multiple hardware streams will hit or miss the first level. See page 47, column 1, and note the paragraph beginning with "Another...". Note that if a miss is predicted at the second level of cache (L2), then a miss is inherently predicted at level one (L1). This is inherent because the nature of cache accessing forces it to be. More specifically, if there is no miss at level one, there

Art Unit: 2183

is no need to access level two. Therefore, in order for a miss to matter at level two, there must have first been a miss at level one.

b) second forecasting whether the corresponding instructions from the corresponding one of the multiple hardware streams will hit or miss the second level. See page 47, column 1, and note the paragraph beginning with "Another..."

c) wherein said selecting comprises choosing the selected one of the multiple hardware streams based upon whether the corresponding instructions from the corresponding one of the multiple hardware streams will hit or miss the second level. Again, see page 47, column 1, and note the paragraph beginning with "Another...". The fetch algorithm includes the steps of fetching from a current thread as long as a miss is not predicted, and when a miss is predicted, switching to and fetching instructions from a second thread.

Claim Rejections - 35 USC § 103

26. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

27. Claim 15 is rejected under 35 U.S.C. 103(a) as being unpatentable over Yoaz, as applied above.

28. Referring to claim 15, Yoaz has taught a system as described in claim 1. Yoaz has further taught that the data cache comprises a first level and a second level (see page 47, column

Art Unit: 2183

1, Fig.3, and the paragraphs beginning with “Employing...” and “Another...”. It is clear that there is an L1 and L2 cache.), and wherein said each of said multiple hit/miss predictors comprises:

a) a hit/miss predictor, configured to forecast whether said corresponding instructions from said corresponding one of the multiple hardware streams will hit or miss said first level. See page 47, column 1, and note the paragraph beginning with “Another...”. Note that if a miss is predicted at the second level of cache (L2), then a miss is inherently predicted at level one (L1). This is inherent because the nature of cache accessing forces it to be. More specifically, if there is no miss at level one, there is no need to access level two. Therefore, in order for a miss to matter at level two, there must have first been a miss at level one.

b) a hit/miss predictor, configured to forecast whether said corresponding instructions from said corresponding one of the multiple hardware streams will hit or miss said second level. See page 47, column 1, and note the paragraph beginning with “Another...”.

c) Yoaz has not explicitly taught separate first and second predictors for predicting first and second level cache misses, respectively. However, as shown in Nerwin v. Erlichman 168 USPQ 177 (1969), to make separable is generally not given patentable weight or would have been an obvious improvement. Doing would allow a user to customize each predictor, separately. Consequently, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Yoaz’s hit/miss predictor into first and second predictors for predicting first and second level cache misses, respectively.

d) wherein said fetch algorithm selects the selected one of the multiple hardware streams based upon whether said corresponding instructions from said corresponding one of the multiple

Art Unit: 2183

hardware streams will hit or miss said second level. Again, see page 47, column 1, and note the paragraph beginning with "Another...". The fetch algorithm includes the steps of fetching from a current thread as long as a miss is not predicted, and when a miss is predicted, switching to and fetching instructions from a second thread.

29. Claims 16-17 and 20 are rejected under 35 U.S.C. 103(a) as being unpatentable over Yoaz, as applied above, in view of Ryan, U.S. Patent No. 5,694,572.

30. Referring to claim 16, Yoaz has taught a system as described in claim 1. Yoaz has further taught a network processor (Yoaz's processor may be use din any way including in a network), but has not taught that said each of said multiple hit/miss predictors employs a flow number to which a packet belongs to forecast whether said corresponding instructions from said corresponding one of the multiple hardware streams will hit or miss the data cache. However, Ryan has taught the concept of operating on first process packets (instructions and data) wherein the cache would be populated with data specifically for the first process. Consequently, when a switch is made to a next process (and group of packets), the cache will not be optimized because it is populated with data from the first process. Therefore, the next process will result in many more cache misses until the cache is populated. See column 2, lines 1-16. As a result, one way a prediction could be made is based on process numbers or flow numbers (note that process numbers are well known and expected in the art as they are used to distinguish processes) because a new flow number would mean that the cache is not optimized for that new flow number and a miss is much more likely. This would be a simple way of making a prediction. As

Art Unit: 2183

a result, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Yoaz to use flow numbers assigned to specific packets for making predictions.

31. Referring to claim 17, Yoaz has taught a processor as described in claim 6. Yoaz has further taught a network processor (Yoaz's processor may be used in any way including in a network), but has not taught that said each of said multiple hit/miss predictors employs a flow number to which a packet belongs to forecast whether said corresponding instructions from said corresponding one of the multiple hardware streams will hit or miss the data cache. However, Ryan has taught the concept of operating on first process packets (instructions and data) wherein the cache would be populated with data specifically for the first process. Consequently, when a switch is made to a next process (and group of packets), the cache will not be optimized because it is populated with data from the first process. Therefore, the next process will result in many more cache misses until the cache is populated. See column 2, lines 1-16. As a result, one way a prediction could be made is based on process numbers or flow numbers (note that process numbers are well known and expected in the art as they are used to distinguish processes) because a new flow number would mean that the cache is not optimized for that new flow number and a miss is much more likely. This would be a simple way of making a prediction. As a result, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Yoaz to use flow numbers assigned to specific packets for making predictions.

32. Referring to claim 20, Yoaz has taught a method as described in claim 11. Yoaz has not taught that said making comprises employing a flow number to which a packet belongs to forecast whether the corresponding instructions from the corresponding one of the multiple

Art Unit: 2183

hardware streams will hit or miss the data cache.. However, Ryan has taught the concept of operating on first process packets (instructions and data) wherein the cache would be populated with data specifically for the first process. Consequently, when a switch is made to a next process (and group of packets), the cache will not be optimized because it is populated with data from the first process. Therefore, the next process will result in many more cache misses until the cache is populated. See column 2, lines 1-16. As a result, one way a prediction could be made is based on process numbers or flow numbers (note that process numbers are well known and expected in the art as they are used to distinguish processes) because a new flow number would mean that the cache is not optimized for that new flow number and a miss is much more likely. This would be a simple way of making a prediction. As a result, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Yoaz to use flow numbers assigned to specific packets for making predictions.

Conclusion

33. Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire **THREE MONTHS** from the mailing date of this action. In the event a first reply is filed within **TWO MONTHS** of the mailing date of this final action and the advisory action is not mailed until after the end of the **THREE-MONTH** shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37

Art Unit: 2183

CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to David J. Huisman whose telephone number is (703) 305-7811. The examiner can normally be reached on Monday-Friday (8:00-4:30).

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Eddie Chan can be reached on (703) 305-9712. The fax phone number for the organization where this application or proceeding is assigned is 703-872-9306.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

DJH
David J. Huisman
October 12, 2004


EDDIE CHAN
SUPERVISOR
TECHNOLOGY
EXAMINER
2100